platform solutions
on-line news for the hardware developer
Issue 6—February 20, 1998

*** SPECIAL INTEL DEVELOPER FORUM ISSUE ***

Feature

PC Industry Growth and Segmentation

Intel CEO Andy Grove describes his view of the segmenting computer market and outlines Intel's strategy for delivering processor and platform products and technologies that enable OEMs to meet the needs of multiple market segments.

Focus

IDF Keynote Addresses

The webcast archive of Andy's entire keynote address from the Intel Developer Forum (IDF) can be seen at the IDF site. Don't miss the other keynote webcasts live:

Thur – Feb 19th, Dr. Albert Gelsinger, Vice President of Business Platforms Group

IDF Coverage – February 17th, 18th, 19th

Check back here each day and you'll find new **Top Stories** and the latest **Tech News** from the Intel Developer Forum in San Jose, California. For an overview of IDF and the stories you can expect to see, read the **Focus** article.

Top Stories

Intel Prepares Developers for Next-Generation Mobile Platform Technologies at IDF

Intel's powerful next-generation Pentium® II processor will pave the way for mobile computers that feature higher performance, manageability and visual computing capabilities, all while keeping system power at a thermally manageable level. At IDF, Intel outlines important design considerations facing the mobile industry.

Introducing the Intel740 Graphics Accelerator
Delivering Arcade-Quality 3D to Volume PCs

Enabling the adoption of balanced, next-generation 3D performance, the Intel740 Graphics Accelerator delivers breakthrough graphics quality. Learn how the Intel740's Hyperpipelined 3D architecture enables new levels of 3D performance. At IDF, many hardware vendors are showcasing Intel740 based solutions.

AGP 4X: Next-Generation Graphics

The building blocks for delivering the Visual Computing initiative have taken hold in the development of a new class of media-rich PCs. As more robust graphics systems and software take shape, new demands are placed on the system architecture. AGP 4X offers 8-times the peak bandwidth of PCI—find out why the move is an important one in the changing graphics landscape.

Top Stories (continued)

$I_2O^{®}$ Technology: Meeting the needs of SHV Server I/O Today

At IDF, Intel outlined the importance of improving the performance of today's and tomorrow's standard high-volume (SHV) server I/O subsystems. Discussions centered on decoupling I/O from the CPU, offloading I/O tasks, and providing additional processing power to scale I/O processing to meet emerging capacity needs. I_2O* has become the catalyst for improving server I/O performance—find out why and how.

IEEE 1394 and USB—Living in Harmony

USB and 1394 will peacefully coexist in tomorrow's PCs, both providing plug-and-play ease of use, but each serving a different class of peripherals. At the Intel Developer Forum (IDF), Intel architects presented the latest implementation details of USB and 1394 on the PC platform.
Learn more about the differences between 1394 and USB from this synopsis.

Platform News and Information

*** Check out our Platforms, Technologies and Events pages ***

Every month we cover the latest developments in platform initiatives and technologies. Our "Platforms" pages provide news on the latest trends and initiatives for the business, home, mobile, server and workstation platforms. Our "Technologies" pages give you quick and detailed information on the industry status of pecific platform technologies, from the emergence of the Accelerated Graphics Port (AGP) to the latest advances in Intel microprocessors, memory, Audio, USB, 1394, DVD, Power Management, and PC 98. Our "Industry Events" page keeps you up to date of upcoming industry gatherings targeted at the platform and peripheral developer.

Technology News

This department is your source for the hottest technology and product announcements, white papers, design guides, specifications, tools and developer events available to the industry.

Reader Services

If you are new to *Platform Solutions* and would like to receive this companion newsletter to the on-line version, please visit *Platform Solutions* on-line and go the "Subscribe Now" section to register and sign up for delivery. The on-line version provides lots of direct links for quick access to the developer information and news reported in each issue, whether it's on Intel's web site or industry web sites.
Please visit the following URL:

http://developer.intel.com/solutions/

We want you to consider *Platform Solutions* as your personal information resource for the Intel architecture platform. If you can help us make it better, or if you have a comment, question or a specific topic you would like to see covered, we want to hear from you.
Please take the opportunity to send us an email with your specific feedback or request to:
platform_solutions@ccm.sc.intel.com

If you do not want to receive this mailing in the future, please send an email to: platform_solutions@ccm.sc.intel.com with "unsubscribe" in the body of the message.

On behalf of all of us at Platform Solutions, welcome to the future of the PC platform!

Platform Solutions Newsletter - Issue 6, February 20, 1998

platform news for the hardware developer

*** SPECIAL IDF ISSUE ***

Table of Contents

Feature:

PC Industry Growth and Segmentation

By Dr. Andrew S. Grove
Chairman and CEO
Intel Corporation

1997 was another good year of growth in the PC Industry. The biggest driver for this growth is Connected Computing (aka the Internet). If the PC industry continues to grow at its historical rate, there are estimates that we will have a world that will have close to a billion computers. Of those, most will be connected to each other via the world wide network known as the Internet. A world of a billion connected computers offers several new opportunities for new kinds of related computing products as well as increased demand for servers to connect all these computers.

A second driver of growth has been the PC Architecture platform moving into new segments of the computing market. 1998 will see PCs move into more segments of the computing market, as shown in the diagram below.

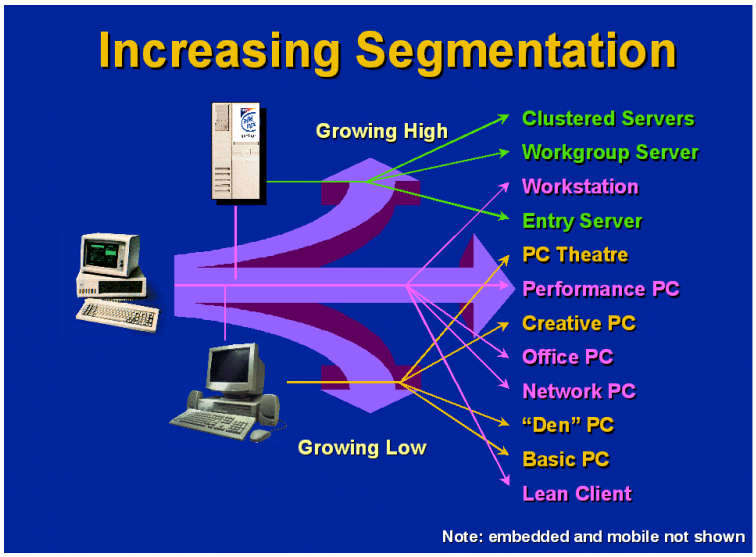

Here we see a wide array of "PC" platforms specifically designed for a particular usage model. Simply put, the PC market is going low and is going high. As a result, Intel is expanding its product portfolio to include processors targeted for all these segments.

Intel will use the P6 Microarchitecture as the underlying foundation for a plethora of new products designed specifically for different segments of computing in addition to the mainstay of our product offering, the Pentium® II processor family of products that are used in the volume of the business and consumer desktop products that are shipped today.

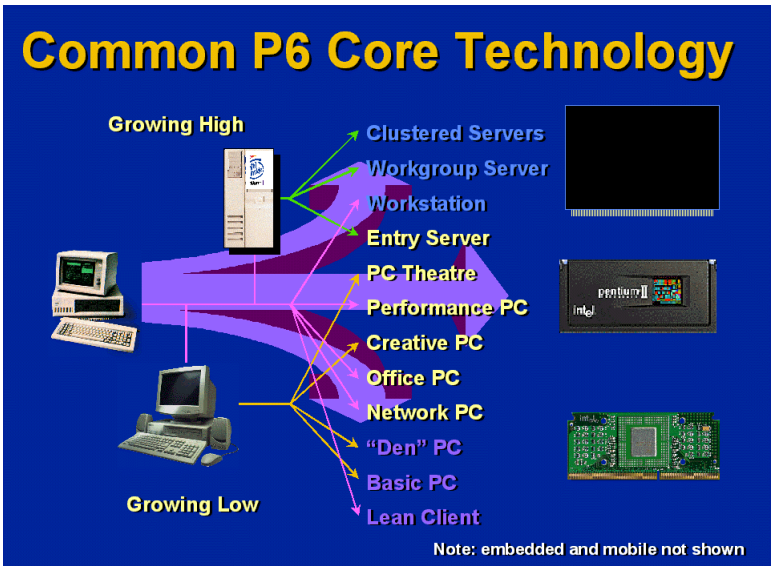

On the low end, we will be introducing a new product for the basic computing design. In order to better serve this segment of computing, we are designing new processors, chipsets and motherboards aimed at low cost designs.

On the other end of the spectrum, we will continue to invest our efforts in expanding our offering to deliver the highest levels of performance for workstations and enterprise class computing. We will be introducing several new products as well as enhanced packaging to allow for faster cache busses and higher frequency processors.

As we have done in the past, Intel will continue to wrap the Intel Inside umbrella around all of our new products. Further we will introduce a new brand name for the basic computing segment, as well as an extension to the Pentium II processor family for the performance segments of computing. The business and consumer desktops, the volume segments of computing today, will continue to evolve and be enhanced under the Pentium II processor family name.

1998 should be an exciting year for both Intel and the computing industry. We look forward to working with you in building products for all the segments of the computing industry.

About the Author
Dr. Andrew S. Grove is the Chairman and CEO of Intel Corporation. He was recently selected as *Time* Magazine's Man of the Year for 1997.

For More Information
To view the webcast of **Dr. Grove's entire keynote address** at the Intel Developer Forum, visit the IDF web site. (http://developer.intel.com/design/idf)

Don't miss the other IDF keynote **webcasts live**:

> **Wednesday – Feb 18th**, Dr. Albert Yu, Vice President of Microprocessor Products Group: Dr. Yu will fully discuss Intel's products targeted at multiple segments.
>
> **Thursday – Feb 19th**, Patrick Gelsinger, Vice President of Business Platforms Group:
> Mr. Gelsinger will discuss Intel's platform and technology efforts and working with the industry to enable development of OEM and IHV products for multiple segments.

Focus:

Intel Developer Forum, February 17th–19th: Three Days that Can Make 1998 a Very Good Year

The new year is only six-weeks old, but 1998 is already shaping as a year filled with profound challenges for the computer industry.

The major challenge: how to meet the growing demands of a highly segmented marketplace for computer products that deliver more performance—within the constraints imposed by volume price points. The challenge may be great, but Intel is now working with the industry to help meet it.

IDF provides "High-Granularity" Information

The second Intel Developer Forum (IDF) is taking place the week of February 17th. Developers who attend this three-day technology-focused event will come away with technical information and "high-granularity" answers they need to raise the price/performance bar to new levels in 1998.

The familiar adage of "faster, better, cheaper" may be a cliché, but it has the ring of truth to most hardware developers today. As the demand for new PCs continues to grow, users in all application segments have learned to demand more performance in volume-priced products. It is a phenomenon that places a set of new design constraints on the industry.

This week at IDF Intel's top architects are showing developers new ways to implement the latest technologies and deliver the products that today's users demand. Baseline specifications are now in place, and new one's are being developed that will enable us to advance all computer platforms through 1998 and beyond. The IDF goes several steps further, offering the detailed technical tracks needed to move from "beyond the specs" to practical implementations.

Three Days Filled with Answers

Tuesday—February 17th

IDF begins on Tuesday, February 17 with a morning keynote presentation by Dr. Andrew Grove, Intel Chairman and CEO. He will present his vision of the changing computer marketplace and Intel's strategic view of the trends, products, and technologies that will shape the evolution of computer platforms in the year ahead.

Tuesday's technical tracks:

- Designing the Basic PC
- High-Performance Memory Implementations
- Workstation Platform Technologies
- Server Platform Technologies (Part I)
- Mobile Platform Technologies
- Lab: Wired for Management Baseline Specification R 1.1

IDF also features over fifty technology demonstrations by Intel and other leading companies at the Tuesday evening Demo Showcase.

Wednesday—February 18th

On Wednesday, Dr. Albert Yu, Senior Vice President and General Manager of the Intel Microprocessor Products Group, will keynote the day with a strategic outline of Intel's future product directions.

Wednesday's technical tracks:

- Delivering Cost-Effective, High-Performance PCs
- AGP-4X: Next-Generation Graphics
- External Interconnect Technologies
- Server Platform Technologies (Part II)
- Lab: Wired for Management Baseline Specification R 1.1
- Lab: Mobile Lab—Making Your Notebook WfM-compliant
- Lab: Mobile Lab—Mobile Power Tools Workshop

Thursday—February 19th

Thursday's sessions begin with a keynote presentation by Patrick Gelsinger, Vice President and General Manager of Intel's Business Platform Group, who will discuss platform technology roadmaps for 1998.

Thursday's technical tracks:

- Host-based Interactive DVD
- Wired for Management Baseline Specification
- Understanding 3D Graphics Performance
- IPEAK Tools for Power Management and Mass Storage
- Instantly Available PC Power Management
- Server Platform Technologies (Part III)
- Intel Product Integration
- IA-64 Architecture Innovation
- Lab: Mobile Lab—Making Your Notebook WfM Compliant
- Lab: Mobile Lab—Mobile Power Tools Workshop

Catch the Flavor of IDF in *Platform Solutions* This Month

This month's *Platform Solutions* provides a taste of what attendees can expect at IDF. Some of Intel's experts involved in developing the IDF technology tracks have authored Top Stories bringing you key aspects of the technologies discussed. By reading these stories you will not only get a flavor of what attendees are getting, but you'll get some key technology information as well. Here's a run-down of the Top Stories you'll find if you check back on the front page of *Platform Solutions* each day of IDF.

Brian Ekiss, Intel's Graphics Marketing Manager, delivers a preview of the new **Intel740 Graphics Accelerator** chip and its key technology components. Jim Nucci, Intel's AGP Marketing Manager, offers a technical perspective on next generation graphics with **AGP-4X**, and Bill Pearson, Intel's 1394 Technical Marketing Engineer, outlines what's ahead in connectivity by comparing **1394 and USB technologies**.

Intel's IA-64 Architecture Marketing Director, Ron Curry, offers a look at the benefits of EPIC technologies coming to you in Intel's first **IA-64 Architecture** Merced™ processor. Intel's MD6 Marketing Director, Richard Dracott, shares the 1998 roadmap details of **IA-32 Processors** as discussed in Dr. Albert Yu's keynote address, and Mitch Shults describes how the SHV Server industry is overcoming **Server I/O** bandwidth challenges today with I_2O® technology.

As PCs become more powerful, developers confront new **System Design** issues requiring new solutions as outlined by Martin Rausch, Intel's Electrical Platform Technologies Manager, in the first of an ongoing series of System Design articles. Brian Johnston, Platform Technical Marketing Engineer, provides an example of the newly announced **Graphics IPEAK tool**, while Frank Spindler, Mobile Marketing Director, describes the hottest **Mobile Platform Technologies**, and Patrict Bohart, Platform Technical Marketing Engineer, provides the details of implementing Dual Mode Power Delivery in the **Instantly Available Power Managed PC**.

Don't Miss the latest Technology Announcements from IDF
Platform Solutions will also provide all the technology announcements coming from IDF each day on the front page **Tech News** section. So stay tuned to get all the latest platform technology information coming from this significant event.

To get even more details on the descriptions of the IDF technical tracks, presenters, and even live/archived webcast of the IDF keynotes, please visit the **IDF web site** (http://developer.intel.com/design/idf).

Top Stories:

Intel Processors in 1998: Meeting the Needs of Multiple Market Segments

by Richard Dracott
Director of Marketing MD6
Intel Corporation, Microprocessor Products Group

Intel is bringing the advanced capability and performance of the P6 microarchitecture* to a variety of unique market segments. First introduced with the Pentium® Pro processor in 1995, Intel has continued to develop the P6 microarchitecture into the Pentium® II processor of today with the scalable performance features of Dual Independent Bus (DIB) architecture and Dynamic Execution. These advanced features of the P6 microarchitecture will continue to evolve, offering an affordable and flexible core design capable of supporting the increasingly rich variety of 32-bit PC application segments.

In 1998, the P6 microarchitecture will power an ever more diverse array of processor types uniquely designed for specific PCs, ranging from highly affordable Basic PCs** to volume-priced Performance PCs designed to run the latest **Visual Computing (**http://developer.intel.com/solutions/archive/issue1/focus.htm) applications for businesses, homes and schools. In addition, the versatile P6 microarchitecture processor design will power mobile systems for the first time, and enable a new generation of high-performance servers and workstations.

Here is just one example. Intel just introduced a new version of the Pentium II processor designed to meet the needs of high-performance desktop PCs for both business and consumers, plus entry-level servers and workstations. This newest member of the Pentium II processor family is a Slot 1 implementation that runs at 333 MHz with a 66-MHz system bus and supports 1- and 2-way processing. Like all the new P6 microarchitecture based processors planned for 1998, this new processor is based on Intel's 0.25 micron process technology, enabling manufacturing economy, lower power consumption, and new levels of processor performance.

New Processor Variations
During 1998, Intel has a variety of processors planned, each combining high performance with design characteristics and pricing optimized for specific platform implementations.

Here is a quick sketch of Intel's 0.25 micron P6 microarchitecture roadmap for 1998:

- Coming in the first half of 1998—new Pentium II processors for Slot 1 for performance and professional PCs and entry-level servers and workstations with a 100 MHz system bus, which will also support new high-speed SDRAM technology.
- Also coming in the first half of 1998—Pentium II processors for mobile computers, available in both mobile module and mini-cartridge form factors.
- Coming in mid-1998—a Slot 1 implementation which will bring the P6 microarchitecture* performance to Basic PCs** at sub-$1,000 price points.
- In the second half of 1998—the introduction of the P6 microarchitecture in the Slot 2 form factor with a 100-MHz system bus and optimized L2 cache configurations designed specifically for mid-range to high-end servers and workstations.

The Need For Speed
In January of 1998, the Pentium II processor reached 333 MHz on the current 66-MHz system bus. With the implementation of the new 100-MHz system bus the Pentium II processor Slot 1 implementation will attain speeds of 350, 400 and 450 MHz, while Slot 2 processors focused at the high-end are expected in 400- and 450-MHz versions. The new Pentium II processor for mobile will enable platforms with speeds of 233, 266 and 300 MHz before the end of 1998. The bottom line is that even as the role of PC platforms continue to diversify and specialize, OEMs will be able to optimize performance for virtually any PC usage model and price point while simultaneously reaping all the benefits of a common P6 core architecture.

Meeting The Needs Of Basic Pcs
Basic PCs in the sub $1,000 price category represent a fast-growing percentage of all PCs sold. Systems in this price-sensitive segment are often purchased as a second PC for homes and very entry-level systems for businesses, as well as attracting first-time computer buyers. Intel's P6 microarchitecture based processor for this segment, code named Covington, is uniquely designed to meet the cost/performance balance for this market segment. Intel will provide this balance by implementing ways to reduce the cost of the processor (e.g. offering a version without L2 cache). Future versions of P6 microarchitecture processors for the Basic PC market segment, expected by the end of 1998, will integrate an on-die L2 cache. Dynamic Execution, the multi-transaction system bus, 32K L1 cache, and all of the other key features of today's P6 microarchitecture, will also be present in these future offerings for the Basic PC.

Slot 2 Processors For Servers And Workstations
Designed specifically for mid-range to high-end servers and workstations, the Slot 2 implementation of the Pentium II processor adds the performance advantage of a full-speed L2 cache bus (L2 cache runs at the full speed of the processor), supporting a variety of L2 cache sizes. L2 cache will be available in 512-Kbyte, 1-Mbyte and 2-Mbyte configurations by the end of 1998.

Making The Transition To 100-Mhz System Bus
As processor speeds increase, other platform building blocks must keep pace to assure a balanced platform with meaningful performance gains. Intel is now working with the industry to enable platforms based on a 100-MHz system bus, including 100-MHz SDRAM and core logic chip sets with features to support these new products. Although high-performance desktops and entry-level servers and workstations will begin to utilize this new bus speed in the first half of 1998, the volume mainstream transition to 100-MHz system bus products is expected in the second half of 1998.

The Roadmap Diversifies In 1998
The Pentium II processor is here today, providing a flexible platform for the development of server, business, and consumer systems. In 1998, the 0.25 micron P6 microarchitecture will diversify and extend to high-performance PCs, Basic PCs, as well as a wide range of workstations, servers, and mobile PCs. The common core architecture, flexible configurations and price points will also support new emerging platform types like the Network PC (Net PC) in business and family room PCs in the home.

By integrating a variety of L2 cache sizes and speeds, processor and bus speeds, and product form factors, the performance of the P6 microarchitecture will be extended to a new generation of powerful PCs and applications, serving the needs of every customer segment at all price points through 1998 and beyond.

About the Author:
Richard Dracott is Director of Marketing in Intel's Microprocessor Products Group MD6 division. He is responsible for product and technical marketing strategy and communications of Intel's Pentium II processor product line.

For More Information:
To stay on top of announcements and information available to the industry on Intel's microprocessors in 1998, visit the **Microprocessor technology page** in Platform Solutions on a monthly basis. http://developer.intel.com/solutions/tech/micro.htm

For information on Intel's **IA-64™ microprocessor architecture**, (article to be available on 2/19/98) read the other Top Story by Ron Curry, Intel's IA-64 Architecture Director of Marketing, based on the technical track at IDF in this month's special IDF issue of PSN.

*Microarchitecture is the logical design of the processor core and techniques that determine how a microprocessor executes its instructions.
**This data is to be used for industry, analyst, and press communication only. This wording is not to be used in merchandising or other market materials.

Introducing the Intel740 Graphics Accelerator Delivering Arcade-Quality 3D to Volume PCs

By Brian Ekiss
Graphics Marketing Manager
Intel Corporation, Graphics Components Division

Intel is introducing the Intel740 graphics accelerator chip to deliver advanced-game arcade, 3D authoring and digital video technology to affordable PCs. The Intel740 will accelerate the industry's adoption of balanced, next-generation 3D performance on volume PC platforms. It represents the building block that system architects need to fully realize the graphics performance enabled by the convergence of the **Pentium® II processor with Dual Independent Bus (D.I.B.) architecture and Accelerated Graphics Port (AGP)** technology implemented in the Intel 440LX AGPset (http://developer.intel.com/solutions/archive/issue4/stories/top1.htm).

The Intel740 graphics accelerator is the next step toward realization of Intel's **Visual Computing initiative** (http://developer.intel.com/solutions/archive/issue1/focus.htm), an industry blueprint to unleash the visual power of the PC platform. The Intel740 graphics accelerator supports key Visual Computing technologies, including AGP-2X, 3D graphics and DVD video. It also takes full advantage of the performance enhancements of the latest 333 MHz Pentium II processor and sets the stage for future Intel AGPsets.

Enhancing AGP performance
Intel's 440LX AGPset enhances system bandwidth and concurrency by means of a feature called "Quad Port Acceleration" (QPA). In essence, QPA allows the AGPset to concurrently arbitrate the processor bus, the graphics bus, the PCI bus and SDRAM. By providing faster data transfer rates to the graphics accelerator, AGP eliminates the bottleneck which would exist if the system relied on the PCI bus alone for the transfer of graphics data. The AGP interface makes it possible for software developers to create more realistic 3D applications with the use of large detailed texture maps that can be stored and executed directly from system memory, thus avoiding the constraints of 2MB or 4MB local graphics memory.

The Intel740 graphics accelerator is an AGP-2X device. Typically the graphics accelerator sends data to the CPU once at every peak of the clock cycle. In an AGP-2X device, the graphics accelerator sends data to the CPU at both the rising and falling edges of the clock cycle, which effectively doubles the quantity of information that can be sent in a given time interval.

Balanced performance for DVD playback
The Intel740 is designed to produce the most realistic 3D performance available from the advanced floating-point capability of the Pentium II processor with Dual Independent Bus (D.I.B.) architecture. In addition to enhancing 3D performance, the Intel740 chip also includes a parallel 2D engine optimized for responsive 2D performance, and smooth playback of both hardware-based and software-based (host-based) DVD video.

"Hyperpipelined 3D"
The Intel740 graphics accelerator assures consistent 3D acceleration and the best combination of quality and performance with a unique feature-set known as "Hyperpipelined 3D." The "Hyperpipelined 3D" architecture has three key elements:

1. **Precise Pixel Interpolation (P.P.I.)—**3D images are made up of polygons (triangles). Traditional 3D accelerators perform per-vertex calculations and then copy the results over the entire area of a given polygon, resulting in a blocky appearance and other visual anomalies. By contrast, the Intel740 texture engine dramatically improves realism and delivers greater precision in its interpolation operations by using floating-point sub-pixel and sub-color calculations. Texture and color values are applied to sub-pixel areas within each polygon, resulting in much more lifelike and "highly detailed" 3D effects. Precise Pixel Interpolation means that 3D textures and attributes such as fog and specular effects are rendered in a much more realistic manner.
2. **Parallel Data Processing (P.D.P.) —**The Intel740 graphics accelerator supports concurrent parallel execution of graphics data on the AGP bus—and on the memory bus. At the same time, the Intel740 supports parallelism inside the chip, enabling the concurrent interpolation of up to 15 different 3D attributes, including atmospheric, perspective, color, specular, translucency and depth data. This parallel processing capability allows several commands to be queued at the same time in the graphics pipeline, translating into consistently lifelike 3D performance, regardless of the complexity of a scene. According to John Carmack, technical director and president of ID Software, "The Intel740 3D accelerator provides no-compromise performance for Quake II players. Intel has shown that it isn't necessary to cut features or sacrifice quality in consumer level 3D."
3. **Direct Memory Execution (D.M.E.)—**By making use of the Quad Port 440LX AGPset, the Intel740's deeply buffered 3D pipeline can execute large, detailed texture maps while they reside in system memory. The advantage of doing it this way is that valuable bandwidth is saved and much larger textures—as large as 30MB and beyond—can be used to create stunning 3D applications.

Intel's commitment to the industry

With the Intel740 graphics accelerator, Intel is enabling the industry to deliver balanced, next-generation 3D performance on volume PC platforms. When combined with the Pentium II processor and 440LX AGPset, the Intel740 graphics accelerator delivers compelling 3D and visually realistic experiences that PC users want and expect on high-performance mainstream PCs in 1998.

About the Author

Brian Ekiss is Graphics Marketing Manager of Intel's Graphics Components Division. He is responsible for the planning and delivery of high-performance graphics building blocks to enable visual computing on volume PC platforms.

For More Information

Visit the new **Intel740 graphics accelerator web site** for more product details (http://developer.intel.com/design/graphics/intel740.htm).

For more information on 3D graphics technology visit Intel's new **3D graphics web site** (http://developer.intel.com/technology/3d).

The Intel740 chip is also integrated into complete solutions available from a **number of leading manufacturers** of AGP graphics accelerator products including:

• ASUStek	http://www.asus.com.tw	• Number Nine	http://www.nine.com
• Diamond Multimedia	http://www.diamondmm.com	• Real3D	http://www.real3D.com
• Leadtek	http://www.leadtek.com	• STB	http://www.stb.com

For more information on **AGP and DIB architecture**, read the top story in Issue 4 of *Platform Solutions* by Richard Malinowksi, Intel's Director of Chipset Engineering

http://developer.intel.com/solutions/archive/issue4/stories/top1.htm).

Intel Prepares Developers for Next Generation Mobile Platform Technologies at IDF

By Frank Spindler
Marketing Director
Intel Corporation, Mobile Handheld Products Group

Intel's powerful next generation Pentium® II processor will pave the way for mobile computers that feature higher performance, are more manageable and provide a wider range of advanced visual computing capabilities, all while keeping system power at a thermally manageable level.

As next-generation mobile computing systems based on the Pentium II processor begin to emerge on the product development drawing board, performance, power consumption and manageability are important design considerations for mobile PC OEMs and Independent Hardware Vendors (IHVs). To help address the latest set of challenges and opportunities for the mobile industry, Intel presented a substantial mobile computing track at the February 17th–19th Intel Developer Forum (IDF).

Covering such topics as the Mobile Power Guidelines and thermal design techniques, mobile Wired for Management (WfM) requirements and recommendations for designing power- and thermal-efficient graphics subsystems for the mobile PC platform, IDF provided the tools OEMs and IHVs need to implement next generation technologies.

Mobile Power Initiative
Thanks to architectural enhancements, new process technology and lower voltage levels, the mobile Pentium II processor has been designed to stay within the targets recommended in Intel's Mobile Power Guidelines. One of the big highlights of IDF was a technology demonstration during the keynote presentation of Dr. Albert Yu, Intel's Senior Vice President of Microprocessor Products Group. The demonstration featured the as-yet-unannounced 266-MHz mobile Pentium II processor running in a notebook. It was shown consuming about the same power as a 166-MHz mobile Pentium processor with MMX™ technology, while providing the significant performance gains expected of the Pentium II processor generation.

Power dissipation remains one of the primary hurdles to overcome as future generations of mobile computers pack increasing levels of performance and functionality into ever-shrinking system footprints. The Mobile Power Guidelines, announced in September 1997, represent the efforts of Intel and leading mobile computer manufacturers, component suppliers and software vendors to address the many power challenges facing the industry.

At IDF, Intel discussed how developers can design systems to meet power targets set in **version 1.0 of Intel's Mobile Power Guidelines '99** (http://developer.intel.com/design/mobile/intelpower/). As shown in the "Managed vs. Unmanaged Power" chart below, the challenge moving ahead is to find ways to limit maximum peak system power to less than 23-25 watts, in order to provide adequate thermal cooling and extend battery life.

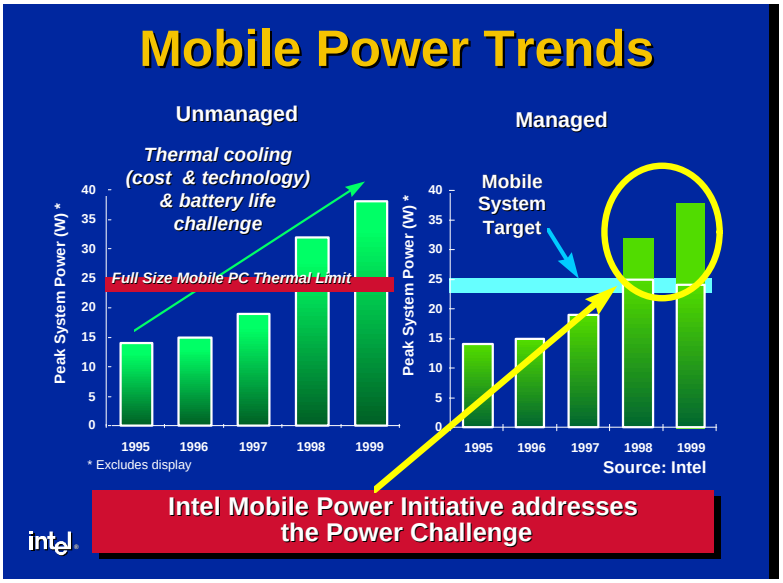

The Mobile Power Tools lab at IDF also featured hands-on demonstrations of Intel's Power Monitor, Power Analyst and Power Management and Analysis tools, which help developers to analyze the power efficiency of various mobile subsystems and applications during design, development and validation. For more information on these tools please visit the Mobile Power Initiative web site (http://developer.intel.com/design/mobile/intelpower/).

Mobile Manageability

IDF discussions also centered around mobile Wired for Management (WfM), a part of Intel's **WfM initiative** (http://developer.intel.com/solutions/tech/wfm.htm) to reduce the total cost of ownership for the enterprise. The industry is adopting a mobile-specific WfM Baseline set of specifications (version 1.1), covering such requirements as component instrumentation, remote wake-up, service boot capability—and of course, power management.

The mobile Manageability Lab sessions at IDF covered installation, operation, debug and customization of Intel's Mobile Component Instrumentation (CI) SDK, optimized specifically for mobile PCs. The CISDK provides a quick and easy way to implement mobile instrumentation that is compliant with the WfM Baseline Specification. Intel also demonstrated its recently announced **LANDesk® Client Manager 3.2** (http://www.intel.com/pressroom/archive/releases/LD021098.HTM), to show how a mobile PC's health can be assessed over a phone or local area network.

Visual Computing Comes to the Mobile Platform

Another important area of focus at IDF centered around the design changes and challenges related to providing more robust visual computing capabilities on the mobile platform. Here also, power is a major consideration, as Intel pointed out by recommending an upper limit of 2.6 watts for the entire graphics subsystems of mobile computers coming to market in the second half of 1998.

Graphics subsystem design challenges discussed at IDF include the transition from PCI to AGP (Accelerated Graphics Port) and the many challenges it raises: increases in platform power and EMI, tighter bus timings and related power management and software driver considerations. Attention was also devoted to host-based DVD, 3D and frame buffer challenges, all driven by end-user demand for ever-higher performance and features at lower price points.

Mobile Opportunities

From power considerations to manageability to the integration of advanced visual computing capabilities, the mobile computing landscape is undergoing a significant change. With the advent of the Pentium II processor, IHVs and OEMs are now in a position to take advantage of a wealth of new opportunities to add their particular value and expertise to the mobile platform. As demonstrated at IDF, Intel continues to work with the industry to expand these opportunities and to provide guidelines that pave the way for higher mobile performance and more features, while keeping power consumption and total cost of ownership down.

About the Author

Frank Spindler is Marketing Director of Intel's Mobile Handheld Products Group (MHPG). He is responsible for marketing Intel's complete line of mobile processors and platforms.

For More Information

Intel **Mobile Power Initiative**:
http://developer.intel.com/design/mobile/intelpower/

Intel **Wired for Management initiative**:
http://www.intel.com/managedpc/

Latest **Mobile platform technology news**:
http://developer.intel.com/solutions/platfms/mobile.htm

IA-64™ Architecture Delivers New Levels of Performance, Scalability

by Ronald E. Curry
IA-64™ Architecture Director of Marketing
Intel Corporation, Microprocessor Products Group

New applications for Internet communications, e-commerce, compute-intensive CAD, 3D graphics and the rapid growth of the **Visual Computing model** (http://developer.intel.com/solutions/archive/issue1/focus.htm) all place increasingly heavy work loads on high-performance workstations and servers. While the performance of today's processors continue to improve, existing architectures based on an out-of-order execution model require increasingly complex hardware mechanisms and are increasingly impeded by performance limiters such as branches and memory latency.

Intel's IA-64™ processor architecture is designed to overcome these limitations. In addition, the IA-64 architecture provides additional performance headroom and scalability needed for future, compute-intensive applications. The IA-64 architecture features a revolutionary 64-bit instruction set architecture (ISA) which applies a new processor architecture technology called EPIC. Jointly defined with Hewlett-Packard Company, EPIC, or Explicitly Parallel Instruction Computing, embodies a set of advanced computer architecture techniques such as *explicit parallelism*, *predication,* and *speculation*. These techniques, as applied to Intel's IA-64 architecture, enable a much higher degree of *instruction-level parallelism* (ILP), and enable IA-64 processors to execute more instructions per clock cycle to deliver superior performance relative to today's out-of-order based RISC processors.

Explicit Parallelism
In today's processor architectures, the compiler creates sequential machine code that attempts to imply parallelism to the hardware. The processor's hardware must then reinterpret this machine code and try to identify opportunities for parallel execution—the key to higher performance. This process is inefficient not only because the hardware doesn't always interpret the compiler's intentions correctly, but also because it uses valuable die area that could be better used to do real work—like executing instructions. Even today's fastest and most efficient processors devote a significant percentage of hardware resources to this task of extracting more parallelism from software code.

The IA-64 architecture's use of *explicit parallelism* enables far more effective parallel execution of software instructions. In the new IA-64 architecture model, the compiler analyzes and *explicitly* identifies parallelism in the software at compile time. This allows the most optimal structuring of the machine code to deliver maximum *Instruction-Level Parallelism* before the processor executes it rather than potentially wasting valuable processor cycles at run time. The result is significantly improved processor utilization. Also, there is no wasting of precious die area for the hardware reorder engine used in out-of-order RISC processors.

Predication Enhances Parallelism
Simple decision structures or *code branches* are a severe performance challenge to out-of-order RISC architectures. In the simple "if-then-else" decision code sequence shown in the figure below, traditional architectures view the code in four basic blocks. In order to continuously feed instructions into the processors instruction pipeline, a technique called *branch prediction* is commonly used, as the name suggests, to predict the correct path. With this technique, mispredicts commonly occur 5 to 10 percent of the time causing the entire pipeline to be purged and the correct path re-loaded. A misprediction rate of just 5 to 10 percent can slow processing speed as much as 30 to 40 percent.

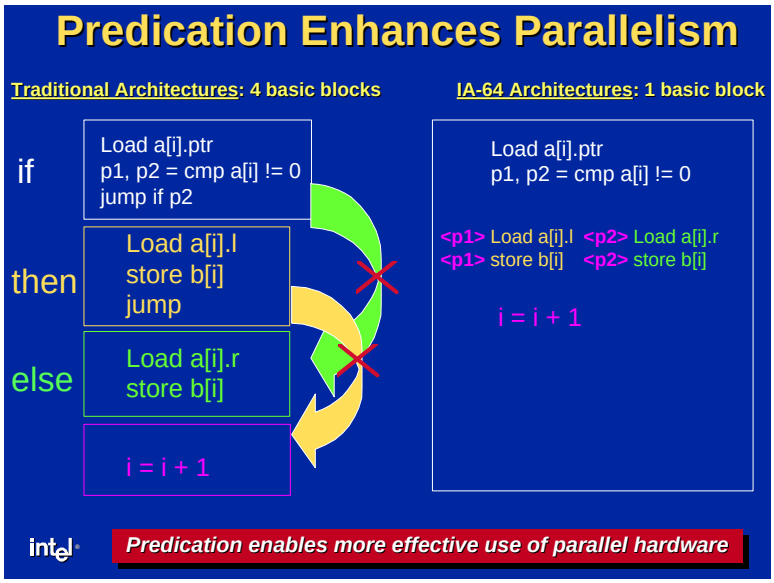

To address this problem and improve performance, the IA-64 architecture uses a technique known as *predication*. In the example above, predication begins by assigning special flags called *predicate registers* to both branch paths—"p1" to the "then" path and "p2" to the "else" path. At run time the compare statement stores either a true or false value in the 1-bit predicate registers. Both paths are then executed by the processor but only the results from the path with a true predicate flag are used. Branches, and the possibility of associated mispredicts are removed, the pipeline remains full, and performance is increased.

Predication is widely applicable. According to a study based on popular software benchmarks (ISCA '95, S. Malhlke et al.), predication can, on average, reduce the number of branches by more than 50 percent and reduce mispredicts by as much as 40 percent. In contrast to some existing architectures, the IA-64 architecture allows all instructions to be predicated.

Speculation Minimizes the Effect of Memory Latency

Overcoming memory latency is another major performance challenge for today's processor architectures. Because memory speed is significantly slower than processor speed, the processor must attempt to load data from memory as early as possible to insure the data is available when needed. Traditional architectures allow compilers and the processor to schedule loads before data is needed, but branches act as barriers to this *load hoisting*.

IA-64 architecture employs a technique known as *speculation* to initiate loads from memory earlier in the instruction stream—even before a branch. Because loads can generate exceptions, a mechanism to ensure that exceptions are properly handled is needed to support speculation that hoists loads before branches.

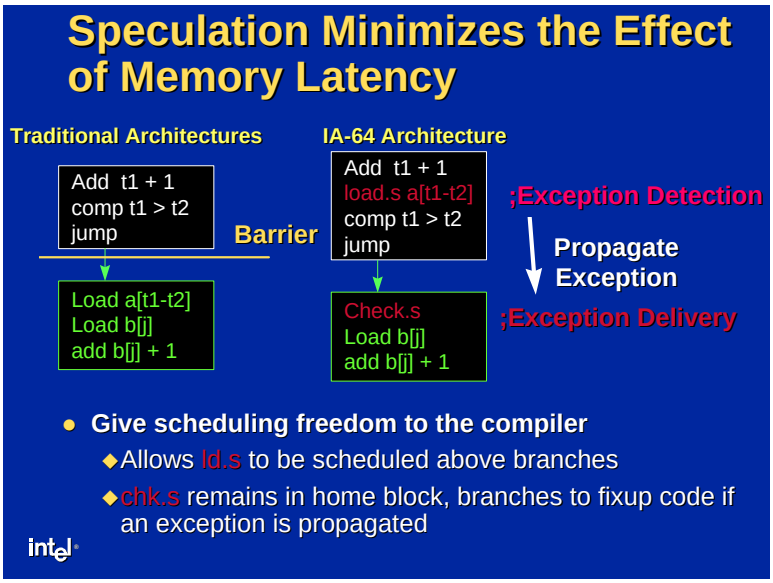

In the above figure, the memory load is speculatively scheduled above the branch in the instruction stream so as to start the memory access as early as possible. If an exception occurs, this event is stored and the *check.s* instruction causes the exception to be processed. The elevation of the load allows more time to account for memory latency, without stalling the processor pipeline.

Branches occur with great frequency in common software code sequences. The unique ability of the IA-64 architecture to schedule loads before branches significantly increases the number of loads that can be speculated relative to traditional architectures. On average, over half of all loads can be executed speculatively resulting in significant performance improvement for today's software.

IA-64 Processors are Massively Resourced and Inherently Scalable

IA-64 architecture innovations enable greater parallelism than traditional architectures. In order to realize the performance improvements of this greater parallelism, the processor must provide massive hardware resources. IA-64 processors include 128 general-purpose integer registers, 128 floating-point registers, 64 predicate registers, and many execution units to ensure enough hardware resources for today's demanding software. The IA-64 architecture is also inherently scalable allowing for straightforward expansion of the number of hardware execution units and increased parallel execution in new IA-64 processor implementations. This is only limited by available silicon technology and provides maximum headroom for future scalability.

Performance, Headroom, Binary Compatibility

The IA-64 architecture's features were designed to bring new capability and levels of performance to server and workstation applications. Today, Intel is leading many industry initiatives to increase system bandwidth and provide balanced platform performance to complement the processing power of the IA-64 processor family. IA-64 architecture processors, the first of which will be Intel's Merced™ processor, will not only enable industry-leading performance, they will offer headroom for future performance gains, while maintaining full binary compatibility in hardware with IA-32 software.

About the Author

Ron Curry is Director of Marketing for Intel's Microprocessor Products Group Santa Clara Processor division. He has marketing, technical support, and industry enabling responsibility for Intel's IA-64 architecture products.

For More Information

To stay on top of announcements and information available to the industry on Intel's microprocessors in 1998, visit the **Microprocessor technology page** in *Platform Solutions* on a monthly basis (http://developer.intel.com/solutions/tech/micro.htm).

For more information on Intel's **IA-64 microprocessor architecture** and Intel's server technologies, read the **Focus article on SHV Servers** by Intel Fellow Justin Rattner in Issue #5 of *Platform Solutions* (http://developer.intel.com/solutions/archive/issue5/focus.htm).

AGP 4X: Next Generation Graphics

By James A. Nucci
AGP Marketing Manager
Intel Corporation, Platform Marketing

It has been nearly one year since Intel formally commenced its **Visual Computing Initiative** (http://developer.intel.com/solutions/archive/issue1/focus.htm). The intent was to accelerate the flow of technology from workstation-class platforms to mainstream-priced PCs, with special emphasis on enabling high-performance graphics. Taking a moment to reflect over the past year's accomplishments, one discovers three prominent advances that are making high-performance graphics on a mainstream-priced PC a reality: (1) Intel delivered the Pentium® II processor with Dual Independent Bus (DIB) architecture, which has the necessary "horsepower" to support the high-performance demands of graphically rich applications; (2) Intel delivered the 440LX AGPset which enabled significant performance gains in various platform subsystems, such as graphics and memory; and, (3) many companies across the industry delivered Accelerated Graphics Port (AGP) products in record time. Building on last year's accomplishments, let's examine today's changing graphics landscape and provide a glimpse into the future of next-generation graphics by way of AGP 4X.

The development cycle for today's AGP graphics products occurred at a record pace. The final version of the *Accelerated Graphics Port Interface Specification Rev. 1.0* was published on the Internet in August 1996, and the introduction of platforms that supported AGP occurred in August 1997—one year from final specification to consumer availability. Those graphics IHVs who invested the resources to design new products to take advantage of the benefits of AGP are seeing great acceptance in the market. For these companies, the words "FreeD" are a forgotten memory. In short, those graphics IHVs who saw how AGP would dynamically affect the market were able to successfully capitalize on this opportunity by investing in the development of products that had the right features and performance to support Pentium II processor platforms. This in turn has caused the landscape of successful graphics companies to change significantly and a new set of product leaders is emerging. Some graphics IHVs still need to rise to the challenge and design in the high-performance AGP features such as AGP Texturing and SideBand Addressing. Having a common foundation of high-performance AGP features can only help support the development of compelling content by ISVs.

The next significant development horizon will occur with AGP 4X products. The ***AGP Interface Specification Rev 2.0*** (http://www.intel.com/pc-supp/platform/agfxport/index.htm) was recently published in December 1997 and included support for AGP 4X. OEMs are already beginning early development work on platforms that support AGP 4X and systems are expected to be available to consumers in the first half of 1999. Once again those graphics IHVs that can manage the transition by developing products with the right features and performance to support future Intel CPUs and platform partitioning are likely to have compelling products to support OEM systems.

The most salient feature of AGP 4X is the peak bandwidth increase to 1GB/s, a peak rate that is two times the bandwidth of today's AGP 2X products. This increase in bandwidth will provide an opportunity for IHVs to develop 3D products with performance levels never seen on a mainstream-priced PC and allow ISVs to create 3D applications that will amaze PC users. If one examines the rapid pace at which the graphics subsystem performance infrastructure has been increasing, as shown in the figure below, you can see the increasing importance of high-performance graphics in the mainstream-priced PC.

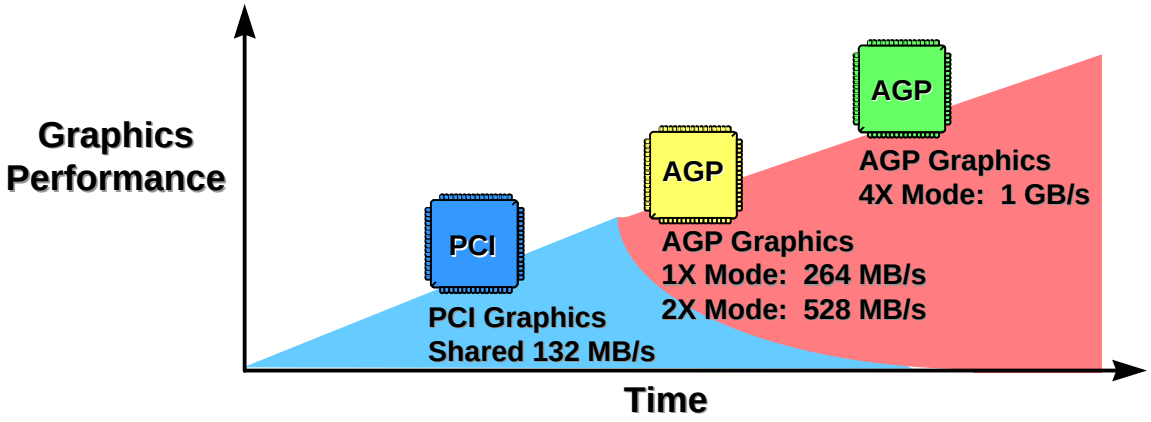

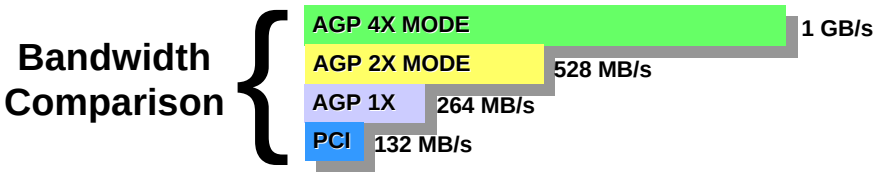

AGP 4X Offers 8-times The Peak Bandwidth Of PCI

Since AGP 4X uses the same underlying protocol as today's AGP, no graphics IHV is likely to have a better understanding of AGP than any other IHV. As I see it, the real challenges with AGP 4X will be designing products that: (1) support the new voltage level required by the high-speed nature of AGP 4X; (2) support the high-performance AGP features like Texturing and SideBand addressing mentioned above; (3) offer the right features and performance at a compelling price; and, (4) have product ready for the introduction of new systems by the PC OEMs in the first half of 1999. Most, if not all of these challenges are firmly in control by the graphics IHVs. Although one can never guarantee success, those graphics IHVs that are committed to accepting the aforementioned challenges and willing to invest their resources appropriately are likely to emerge with successful products.

About the author

James A. Nucci is the AGP Marketing Manager in Intel's Platform Marketing group. He is responsible for working with various vendors throughout the industry to develop and promote AGP-enabled products and systems.

For More Information

To stay in touch with the latest news on AGP, visit the **AGP technology page** in *Platform Solutions* often (http://developer.intel.com/solutions/tech/agp.htm).

For more information on 3D graphics technology visit Intel's new **3D graphics web site** (http://developer.intel.com/technology/3d).

For more information on **AGP and DIB architecture**, read the top story in Issue 4 of Platform Solutions by Richard Malinowksi, Intel's Director of Chipset Engineering (http://developer.intel.com/solutions/archive/issue4/stories/top1.htm).

I2O® Technology: Meeting The Needs Of SHV Server I/O Today

By Mitch Shults
Director of Server Platform Marketing
Intel Corporation, Enterprise Server Group

Microprocessor performance continues to improve rapidly, as measured by the current state of the 32-bit Intel Architecture (IA), and the forthcoming 64-bit IA-64™ Architecture. Combined with similar advances in symmetric multiprocessing (SMP) technology, these performance increases are propelling IA standard high-volume (SHV) servers into the highest rungs of enterprise computing.

In order to provide information technology professionals with the throughput and system scalability they require, SHV server I/O will need to keep pace with continued improvements in processing performance. Thanks to a number of significant technology advances now in place or under development, 1998 promises to be an active year for enhancing the I/O capabilities of the SHV server platform.

At the February '98 Intel Developer Forum (IDF), attention was devoted to improving the performance of today's and tomorrow's SHV server I/O subsystems. Discussions centered on measuring and designing higher performing systems with an eye to meeting ever-increasing throughput and high-bandwidth needs. And finally, IDF focused considerable attention on today's emerging SHV server I/O standard—Intelligent Input/Output, commonly referred to as $I_2O^{®}$ Architecture.

Raising the Bar with 64-bit Multiple PCI Buses
In concert with defining a scalable I/O architecture, Intel is working with the industry to develop chip set technology that brings multiple, high-speed PCI buses to the SHV server platform. Intel's founding efforts in driving the PCI bus standard during the early 1990s helped to establish PCI as the predominant SHV server I/O bus standard today. Our goals in 1998 are to move beyond today's prevailing 33-MHz 32-bit PCI architecture. Initial plans call for the introduction of chip set products that deliver more PCI performance and capacity with 64-bit 33-MHz PCI, including multiple 64-bit PCI buses, by mid-year. Intel will continue this trend with future products incorporating 66-MHz 64-bit PCI server I/O technology by the close of 1998.

I_2O Technology—the new Scalable I/O Architecture
Legacy I/O limitations in Intel Architecture servers exist today because the CPU is optimized for memory access instead of I/O. Single-pass operations such as checksum generation lead to a significant amount of "cache thrashing." The impact of context switches to service I/O interrupts often negates the advantages of pipelines and caches. The goal across the industry is to find better and faster ways to bring data to the CPU—rather than requiring the processor to "slow down" and go out across the bus and fetch the data itself.

I_2O technology defines *the* server I/O specification of today and tomorrow, for the first time providing the industry with a path to achieve I/O performance approaching that achieved by leading mainframe architectures. Intended to offload low-level I/O processing to special I/O processors, the I_2O architecture specifies a split driver architecture for various classes of I/O devices, including block storage (SCSI, tape, etc.) and LAN. Initial tests indicate that I_2O-based adapters should help to significantly offload I/O tasks from host CPUs, and thus substantially increase the scalability of server I/O subsystems. The I_2O architecture decouples the I/O subsystem from the CPU, the OS and even the bus technology, allowing innovations to proceed unencumbered by other system technology restraints.

There were 12 leading server companies that announced I_2O-based server availability at Networld + Interop in October 1997. The first among a number of I_2O-based storage adapters is coming to market shortly, and I_2O-based LAN solutions are also expected in 1998. The next steps in the I_2O specification are also well underway with the next revision adding architectural support for remote I/O adapters attached to the host via a system area network (SAN). Other new features such as peer-to-peer data transfer will also be included.

Next Generation—I_2O and the VI Architecture

Promoted by Intel, Compaq* and Microsoft*, and supported by other leading server vendors, the new Virtual Interface (VI) Architecture is an open specification for low-overhead message-passing over SANs with clusters of volume servers and workstations. Using the VI Architecture, clusters of IA-based SHV servers can be used to address scalability and availability solutions at all levels of the enterprise. As I_2O technology and the VI Architecture evolve on the SHV server platform, together they will form a combined I/O architecture with the scalability needed to fully support the growth of IA processing performance well into the future.

A Time for Action

With all these developments taking place on the server I/O landscape today, the time has come for OEMs and IHVs to focus on developing and delivering I_2O technology solutions for SHV servers to have scalable I/O subsystems that will easily adapt to meet future needs. In addition, OEMs and IHVs need to provide 64-bit multiple PCI bus implementations, dual-address cycles for high-throughput applications and dual-mode subsystem solutions capable of addressing either
33- or 66-MHz implementations.

By any estimation, SHV server I/O technology is on the road to higher performance, scalability and functionality—and Intel is doing its part to help pave the way.

About the Author

Mitch Shults is Director of Server Platform Marketing in Intel's Enterprise Server Group. He is responsible for driving SHV server platform technology strategies and adoption within Intel and the server industry.

For More Information

For more details and news on I_2O technology, visit the **I_2O Technology** page in *Platform Solutions* updated on a monthly basis (http://developer.intel.com/solutions/tech/i20.htm).

For more details and news on VI Architecture, visit the **VI Architecture Technology** page in *Platform Solutions* updated on a monthly basis (http://developer.intel.com/solutions/tech/via.htm).

For information and news on other SHV server platform technologies, please visit the **Server Platform** page in *Platform Solutions* (http://developer.intel.com/solutions/platfms/server.htm).

Be sure to **check out the Server focus Issue #5** of *Platform Solutions* in the Archive with lots of top stories on the hottest SHV server technologies, including a feature article from Intel's Vice President of the Enterprise Server Group, John Miner, and a detailed focus article from Intel's Server Architecture Lab with Intel Fellow Justin Rattner and Director Paul Close (http://developer.intel.com/solutions/archive/).

USB and 1394–Living Together in Harmony

By Bill Pearson
Technical Marketing Engineer
Intel Corporation, Platform Marketing

In the December issue of *Platform Solutions*, we asked the question "Will 1394 replace USB?" The answer, of course, was "No." USB and 1394 are complimentary technologies. USB and 1394 will peacefully coexist in tomorrow's PCs, both providing plug-and-play ease of use, but each serving a different class of peripherals.

At the Intel Developer Forum (IDF) this week, Intel architects are providing a day-long training track discussing the latest implementation details of USB and 1394 on the PC platform. With USB, critical implementation areas are ease of use, bandwidth sharing and power management, as well as platform level interactions between USB and other subsystems. Building on the USB experience, key 1394 requirements in hardware and software building blocks, tools and interoperability are being described, as well as a detailed analysis of the trade-offs of integrating P1394a versus P1394b.

Since they are going to live together, people often ask, what are the differences between 1394 and USB? At IDF developers are getting a detailed look at the differences between USB and 1394. Let's take a look at some of them here.

Topology
The topology of 1394, known as a tree topology, is shown in figure 1 below. Any device can be connected to any other device, so long as there are no loops. A 1394 network can support up to 63 devices. The devices can be hot swapped. If a device is added or removed, the bus will reset, reconfigure, and continue operation. If the bus is broken, the two pieces will reset, reconfigure, and resume operation as two independent busses. 1394 also offers peer-to-peer connectivity, so peripherals can talk to one another without intervention from the PC.

Figure 1—1394 Topology

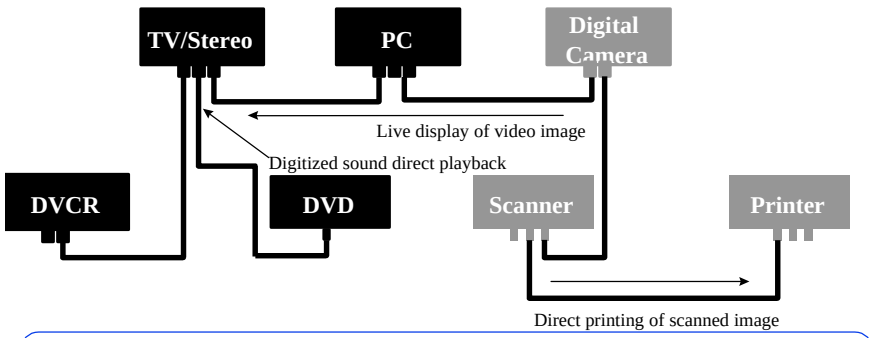

- ***Any node can be a root, typically PC is a root.***
- ***When PC is removed, the separate Serial Busses function normally.***
- ***When PC is reconnected, the two Serial Busses would be reconfigured as one after bus reset.***
- ***No loops allowed.***

In contrast, USB has what is known as a star-tiered topology shown in figure 2 below. The PC acts as the host. Each device is connected to a hub, which provides sockets and power and acts as a repeater. Hubs can be either self-powered or bus-powered. They can also be cascaded. The USB topology supports up to 127 devices.

Figure 2—USB Topology

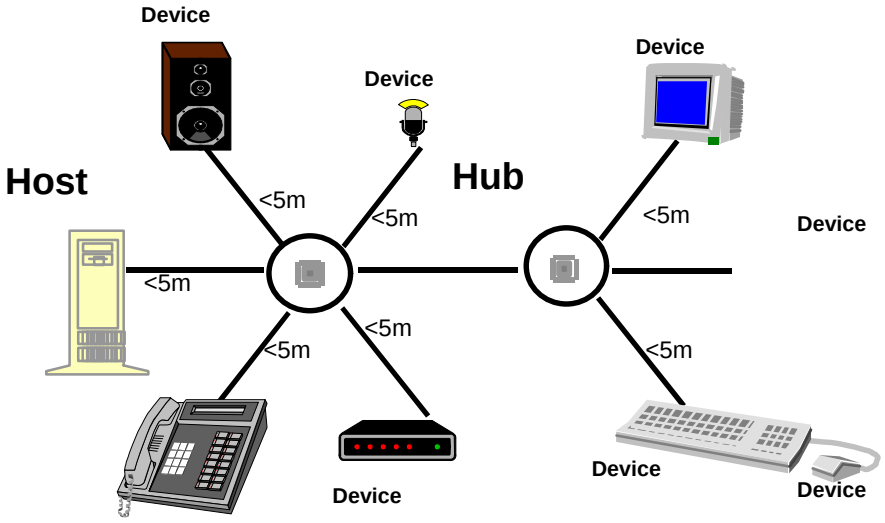

Speed
USB offers speeds ranging from 1 Mbits per second to 12 Mbits per second. In contrast, the current IEEE specification 1394-1995 offers speeds *starting* at 100 Mbits per second and going up to 400 Mbits per second. P1394b will start at 800 Mbits per second and is defining speeds of up to 3200 Mbits per second. P1394b is expected to be fully backward-compatible with the 100-400 Mbits per second specification–connector, cable and software.

Application
As mentioned previously, 1394 and USB are complimentary technologies. USB is a medium bandwidth connection for telephony products, digital still cameras, monitors, keyboards, mice, and other similar I/O devices. In contrast, 1394 is a high-speed bus designed for digital video cameras, DVD players, mass storage devices, and other peripherals that require greater bandwidth.

Cost
USB is a very low-cost interconnect technology. Low-speed USB implementations for devices such as mice and keyboards typically cost less than $1 in OEM quantities, and even the medium-speed implementations for devices like scanners and modems are in the $1-2 range in OEM quantities. Due to relatively lower volumes and higher complexity, 1394 implementations are currently in the $15 range. This cost is expected to decrease as volume builds over the next few years.

Meeting the Needs of Tomorrow's PC
Figure 3 below summarizes the differences between USB and 1394. Tomorrow's PC will have a need for low-cost, low-bandwidth devices like mice, keyboards, and modems. USB provides a good fit for these devices. Tomorrow's PC will also need a high-speed interface for connecting to high-speed printers, hard drives, and camcorders. 1394 is the technology that will allow these devices to interoperate with the PC. Both USB and 1394 will coexist on future PC platforms to meet a wide range of growing peripheral interconnectivity needs.

Figure 3—Comparison Summary

	USB	1394
Speed	Medium Speed 1 Mbit - 12 Mbit	High Speed 100Mbit - 3.2 Gbit
Host Implementation Cost	Low Cost <$1	Moderate Cost ~$15
Topology	Hub with PC in Center	Unsupervised Peer-to-Peer
Applications	Keyboard, Mouse, Joystick, Monitor Hub, Audio, Scanner, Still Camera	D-VCR, DSS, DVD, High Speed Camera, Camcorder, Storage, Printer, Scanner, Device Bay, Networking
Attributes	Ease-of-Use, Hot Pluggable, Higher Capability Peripherals	Ease-of-Use, Hot Pluggable, Digtial Convergence, New Host-Based Applications

Opportunities for OEMs, IHVs and ISVs

You can find both USB and 1394 peripherals in stores today. USB keyboards, mice, and tethered cameras are all available today. There is a significant opportunity for USB peripherals since most PC's are shipping with USB ports today.

The main 1394 peripheral today is the digital, or DV, camcorder. DV camcorders can be found on store shelves today, but are still quite pricey. In 1998, look for 1394 add-in cards on high-performance PCs geared for digital video editing. There is currently a need for 1394 software products for digital video editing and other digital imaging applications, especially products that will work under Windows* 98. In the near future 1394 hard disk drives and DVD drives will begin to appear as well.

Beyond 1998, you can expect to see a higher level of integration with 1394 on the PC platform. As the installed base of 1394-enabled PCs grows, there will be considerable opportunity for 1394 peripherals such as printers and scanners. Intel is currently designing with P1394a rather than IEEE Standard 1394-1995 due to P1394a's significant performance, compatibility and interoperability benefits. P1394a is near its final state and should be submitted to the IEEE balloting process in 1998.

OEMs and IHVs should look today for opportunities to design in support for both USB and 1394 as they will coexist together on the PC platform for the foreseeable future. Look for the next Intel Developers Forum in the second half of 1998 to learn more about the latest USB and 1394 implementation techniques.

About the Author

Bill Pearson is a Technical Marketing Engineer in Intel's Platform Marketing group. He is responsible for industry enabling of 1394 and USB on the PC platform.

For More Information

For the latest news on 1394 including white papers, TA presentations, and specifications, please visit the **1394 technology page** in *Platform Solutions* (http://developer.intel.com/solutions/tech/1394.htm).

For the latest news on USB, please visit the **USB technology page** in *Platform Solutions* (http://developer.intel.com/solutions/tech/usb.htm).

Read the **Q&A on 1394** with Intel's 1394 Program Manager and Technical Marketing Engineer in Issue 4 of *Platform Solutions* (http://developer.intel.com/solutions/archive/issue4/stories/top5.htm).

System Design Technologies Remove Roadblocks to Performance

by Martin Rausch
Manager, Electrical Platform Technologies
Intel Architecture Lab, Platform Technologies Group

The dramatic evolution of processor and PC system performance has raised user expectations to unprecedented levels. Faster processors, memory, and system buses bring with them a host of system design challenges, including electromagnetic interference (EMI), power delivery, interconnect, printed circuit board (PCB) and thermal management design issues.

Because PC users are demanding higher performance at lower price points, the principal design challenge is not simply delivering performance—but delivering this performance at the lowest possible price.

Intel: removing the roadblocks
To meet this challenge, the Platform Technologies Group in the Intel Architecture Lab is chartered with identifying and removing system performance barriers. This group delivers specifications to help vendors develop needed ingredients (connectors, regulators, heatsinks, etc.) and design guides to help OEMs assemble them in high-performance products. At the February Intel Developer Forum platform developers were given detailed implementation techniques in key system design areas. Here is an overview of some of those techniques that shows how Intel is helping the industry meet the performance challenge.

EMI: suppression, not containment
The 66-MHz Front Side Bus (FSB) and SDRAM motherboard clocks and 150-MHz L2 cache Back Side Bus (BSB) interface clocks are the primary EMI radiators in Pentium® II processor-based systems. Future systems will include 100-MHz FSB and SDRAM, 800MTs Direct RDRAM, and 266MTs AGP-4X interfaces, as well as 450-MHz Single Edge Contact (S.E.C.) cartridges housing future Pentium II processors.

The FCC's new open box regulations make containing EMI radiation inadequate. Instead, EMI must be suppressed. Intel is pursuing three technologies to do this with minimal cost impact:

- Complimentary clocks, on all future Pentium II S.E.C. cartridges and the Direct RDRAM interface, reduce emissions by canceling radiation from the primary clock.
- Spread-spectrum clocking, standard on future clock drivers, varies the clock period in small increments to reduce the energy radiated at any given frequency.
- SECC shields inside existing S.E.C. cartridge covers can also be used.

Power delivery solutions
Scaling the voltage downward reduces power but dramatically reduces noise tolerance. Since power isn't decreasing, voltage scaling also creates much higher currents. And, thermal cooling limits demand power-management modes where idle units are powered-down and then quickly powered-up again to resume operation. This causes transient currents (di/dt) whenever a unit is activated or deactivated.

Intel's two pronged strategy will accommodate aggressive voltage scaling and higher transient currents without significant cost increases:

- VRM8.2. VID pins allow VRM8.1 solutions to handle the upcoming 2.8V to 2.0V transition. Processor frequencies above 350 MHz will require VRM8.2 to handle higher maximum currents.

- ≤1.0 nH ESL capacitors, with printed circuit board (PCB) improvements, reduce parasitic inductance instead of using more parallel caps. This handles higher transient currents and lower voltage tolerances without increasing cost.

Interconnect technology
System performance enhancements such as AGP (http://developer.intel.com/solutions/tech/agp.htm) drive rapid bus frequency increases. Trace lengths cannot be shortened enough to meet high bus speed requirements using traditional synchronous designs. Source-synchronous solutions overcome this constraint by sending strobe signals along with the data. Direct RDRAM adds pipelining where new signals are launched before the receiver latches existing ones. Parameters such as clock and data trace length matching, and strobe and data output buffer delay matching—which once could be safely ignored—now limit performance.

Intel has delivered AGP-2X, 66-MHz FSB, and 100-MHz SDRAM memory bus routing guidelines and will release 100-MHz FSB and AGP-4X guidelines by mid-1998. OEMs can simplify the implementation challenge by driving their suppliers to provide:

- Tighter PCB tolerances to reduce the timing skew caused by velocity and impedance variations.
- 3D simulation tools to break the route-simulate loop and visually define how parameter variations affect performance.

PCB design
Delivering solutions in 4-layer motherboards is increasingly difficult. Frequency increases make it harder to meet EMI regulations with clock traces on outer layers. The inductance of standard power and ground connections is increasingly inadequate to meet power delivery requirements. Available routing channels are being consumed by the larger pitch routing required to manage increased crosstalk sensitivity.

Intel is extending the usability of 4-layer motherboards by modifying fabrication and assembly technology:
- Via-in-Pad (VIP) increases routing channels under BGA components and reduces inductance, minimizing the number of decoupling caps required.
- Pads-Only Outer Layer (POOL) moves power and ground planes to the outer layers to increase routing channels, cuts PCB cost, and potentially improves EMI attenuation.

Thermal management
Existing PC chassis designs provide sufficient airflow to cool the processor but not the memory, chip set, or graphics subsystems. A new ingredient called Fan Ducting (FDR) significantly boosts subsystem cooling without extensive retooling of existing ATX and NLX chassis. FDR distributes airflow equally to all platform subsystems eliminating the need for expensive fansinks and heat spreaders.

Providing the design information you need
Intel Architecture Lab is delivering system design technologies so OEMs can deliver higher-performance products at volume price points. As shown at the Intel Developer Forum, Intel is increasing its efforts to work closely with the hardware development community to ensure quick and affordable technology implementation. Look for more system design technology sessions at the Fall IDF, as well as an ongoing series of system design technology news and stories in Platform Solutions.

About the Author
Martin Rausch is Manager of Electrical Platform Technologies in the Platform Technologies Group of the Intel Architecture Lab. His primary responsibilities are developing platform design technologies and techniques that enable system designers to implement new platform technologies at effective costs.

For More Information
You can visit the **Pentium II processor Developer Web** site for lots of system design papers and documentation (http://developer.intel.com/design/PentiumII/).

For the latest news on all the key platform technologies, visit the Technologies pages within ***Platform Solutions*** on a regular basis (http://developer.intel.com/solutions).

New IPEAK Tools Introduced at IDF: New Graphics Performance Toolkit Leads the Way

by Brian Johnston
Senior Technical Marketing Engineer
Intel Corporation, Platform Marketing

Intel's Performance Evaluation and Analysis Kit (IPEAK), originally previewed at last October's Intel Developer Forum, is available now and will begin shipping in March. Developed to help ease technology adoption issues for PC OEMs and Independent Hardware Vendors (IHVs), the IPEAK tools (http://www.intel.com/pressroom/archive/releases/cn021898.HTM) (significantly improve the analytical and diagnostic capabilities inherent in platform integration and performance tuning for the next generation of personal computers.

The first set of IPEAK tools, released after several rounds of beta testing and input from more than 500 reviewers across the industry, is composed of four toolkits designed to improve storage, power management and graphics performance on the PC platform:

- The *IPEAK Storage Toolkit* provides high- and low-level analysis capabilities to optimize the performance of peripherals that affect storage and overall performance.
- The *Intel Power Management Analysis Tool (*IPMAT) provides automated testing capabilities that can save IHVs and OEMs time and money during the product validation and integration process.
- The *Intel Baseline AGP System Evaluation Suite* (IBASES) validates and tests the driver, memory and hardware controller to ensure the delivery of the best AGP solutions.

The IPEAK Graphics Performance Toolkit

The fourth member of the IPEAK family—Intel's *Graphics Performance Toolkit*—enables OEMs, IHVs and Independent Software Vendors (ISVs) to analyze the performance characteristics of graphic controllers and applications, ultimately helping them to optimize the performance of their hardware and software products. The toolkit's second principal benefit is that it can be used to analyze the content characteristics of advanced graphics applications. On the one hand, the *Graphics Performance Toolkit* (GPT) enables IHVs to develop hardware to match the needs of the next generation of graphics applications; on the other, it allows ISVs to create applications that take full advantage of the latest graphics hardware technology.

The *Graphics Performance Toolkit* is unique in that it operates directly on the actual graphics applications being tested. For the first time, OEMs, IHVs and ISVs no longer need to measure synthetic workloads to tune their graphics subsystems or software. Armed with this advantage, the Intel toolkit measures the three primary indicators used today to determine 3D graphics performance: frame rate, triangles and pixels—all measured as a function of time (typically performance per second).

A New Way to Measure True Graphics Performance

In essence, the Graphics Performance Toolkit "snoops" calls to such popular PC graphics APIs as DirectX* and OpenGL*. It times the various routines, collects information, and logs and graphically charts that information over time for use by the reviewer. For example, a developer can launch a 3D application from the GPT and obtain instantaneous frame rate charting of the application's throughput. This enables the developer to gain valuable insight on the processing flow of the application to assist in debug.

The 3D Graphics Pipeline

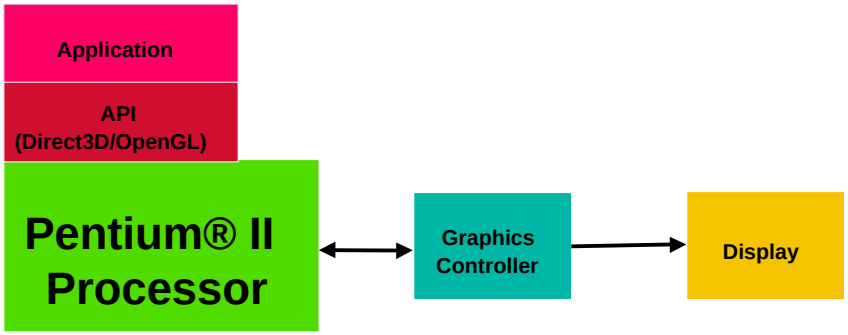

The frame rate measured in frames per second (fps) is the basic metric of graphics software performance. Most, if not all, applications provide the ability to calculate a frame rate mean averaged over a number of frames. While this was sufficient for software rendering, this mean measurement is too coarse for measuring the performance of an accelerated application. Additionally, this "whole pipeline" frame rate can also be subject to interference in the rendering pipeline.

In order to calculate the finest possible measurement, a per frame or instantaneous frame rate (iFPS) needs to be calculated. iFPS reflects the actual amount of time required to handoff a frame of workload from the application to the accelerator.

Graphics Performance Tool (GPT) Screenshot Showing Instantaneous Frame Rate (iFPS)

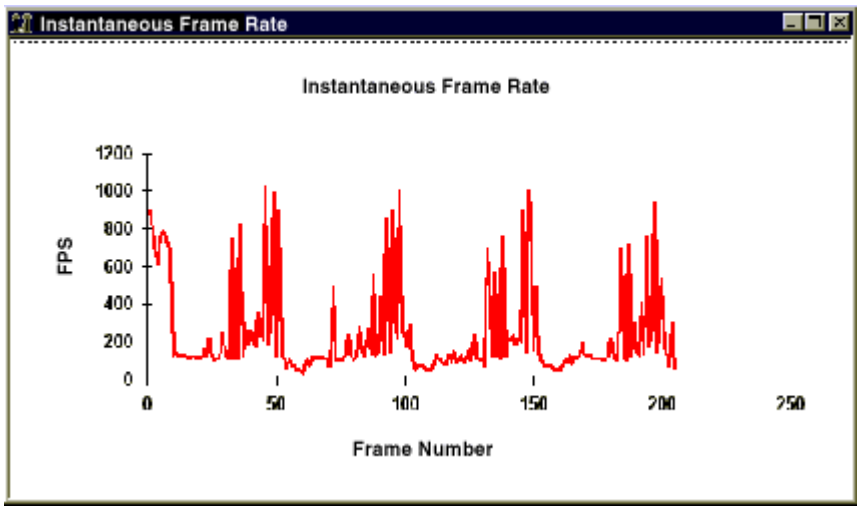

GPT provides the ability to calculate the iFPS as well as to deconstruct the accelerated pipeline by providing full API coverage and instrumentation against the DirectDraw*, Direct3D* IM, Direct3D RM, and OpenGL APIs. This ability allows the GPT to intercept, log, modify, or deflect API functions called from an application in order to more fully understand the performance of the application and its workload.

For OEMs, the *Graphics Performance Toolkit* is more than just a tool to tune their graphics subsystems. By working directly with real applications, the toolkit gives OEMs a way to help determine which software they want to include in their system bundles, and how to make that software run with optimal performance. Similarly, ISVs can enhance the value and attractiveness of their applications by using the *Graphics Performance Toolkit* to tune and improve performance.

To keep up to date on Intel's Platform Performance Tuning technology, visit the **Platform Performance Tuning technology page** in *PSN* often (http://developer.intel.com/solutions/tech/ipeak.htm).

Breaking New Ground

As the first offerings of their kind in the industry, the IPEAK tools help shorten product time-to-market cycles when adopting new platform technologies and standards. As with all IPEAK offerings, the *Graphics Performance Toolkit* provides a way for OEMs, IHVs and ISVs to work together to truly maximize the potential of their Intel Architecture-based hardware and software solutions. The ultimate beneficiaries of these relationships, of course, are end users—who can look forward to enjoying the best possible graphically rich computing experiences available with today's leading-edge PC platform technology.

About the Author

Brian Johnston is a senior technical marketing engineer in Intel's Platform Marketing group. He focuses on developing technical collateral to help accelerate the adoption of platform technologies by the PC industry.

For More Information

For more details on the *Graphics Performance Toolkit* and other Intel IPEAK offerings, please visit the **IPEAK web site** (http://developer.intel.com/design/ipeak).

Implementing Dual Mode Power Delivery on the Instantly Available PC

By Patrick J. Bohart
Technical Marketing Engineer
Intel Corporation, Platform Marketing

A key feature of Intel's Instantly Available Power Managed PC is its support for the S3, or Suspend to RAM (STR), sleeping state. In this S3 state, self refreshing DRAM is used to save and restore system context allowing a noiseless, low power, sleep with the ability to resume to the last system state in only seconds. It is this quick transition from the S3 sleeping state that allows the Instantly Available Power Managed PC to remain connected to the LAN and responsive to an incoming modem call even when "off".

A crucial ingredient necessary in achieving this quick sleeping state transition is support for dual mode output voltages. Dual mode output voltages enable a split power plane design, where the entire motherboard is powered in the active state, while power is only delivered to essential memory, wake devices, and wake logic while in the sleeping (S3) state. Support of dual mode output voltages allow the Instantly Available PC to transition in and out of the sleeping (standby) state cleanly and efficiently.

Intel architects have explored the implementation of dual mode voltages (Vdual) using two different methods. Both methods require the inclusion of a switching mechanism for transitioning between main and standby voltages, yet the two methods vary in the location of this mechanism. One method involves the inclusion of the switching mechanism as a separate circuit incorporated on the motherboard, and the other involves integrating this switching circuitry within the existing circuitry of the power supply.

Options for Implementing Dual Mode Power Delivery

The circuitry involved to support dual mode power delivery in both the motherboard and the power supply implementation consist basically of two FETS; a high capacity FET to route the high runtime current and a p-channel FET to route standby current. In the motherboard implementation this FET circuitry is mated with a standard NLX and ATX power supply. The NLX and ATX power supply provides standard main and standby voltages and the circuitry on the board combines these in a 2:1 mux that provides the Vdual output. Given the 700 mA of standby current available in a standard NLX and ATX supply this configuration is capable of supporting one PCI wakeup device (e.g. LAN card or Modem) once the current requirements of the memory and wakeup logic are subtracted. This implementation is typically viewed as the minimum implementation.

In the power supply implementation the FET circuitry is integrated into the existing power supply circuitry. In this configuration the power supply delivers the standard main and standby voltages with a third dedicated pin delivering Vdual. This new power supply would not be limited by the standby current limitations of a NLX and ATX power supply and could be built to support enough standby current (potentially 2 A) to support up to four PCI wake devices in addition to DRAM and wake logic. This implementation is typically referred to as the maximum implementation.

The Tale of the Tape

Pitting the two methods of implementation against one another may seem like comparing apples and oranges. On the surface the power supply implementation may appear as the lobster, while the FETs on motherboard implementation the crab in a can. Where differences in performance and capabilities are obvious, one might assume that cost is the differentiator, allowing each OEM to choose an implementation that meets their cost vs. performance requirements. This, however, is a much too simplified view of the matchup. Several factors point to the implementation of the dual mode circuitry within the power supply as the superior choice, including reliability, cost reduction, and interoperability.

1. Reliability –
Reliability is a key concern for the Instantly Available PC to operate effectively. In the STR state the majority of the motherboard is completely powered down, and is thus very cool thermally. When the FET circuitry is added to this environment a very localized hot spot is created as the FETs route current and dissipate heat. In this configuration problematic situations such as board warping become a serious concern. Additionally, the integrity of the FET circuitry may be diminished with repeated overheating, potentially reducing the life span of the motherboard.

2. Cost –
Moving the circuitry into the power supply not only removes this hot spot, but can also actually reduce costs by allowing power supply designers to utilize necessary resources already existing within the power supply. The need for an additional heat sink to dissipate heat from the motherboard hot spot can be negated by tapping into an existing heat sink within the power supply. Similarly, bundling the FET circuitry with unused circuits that already exist within the power supply unit can save further costs. Current and voltage overprotection can also be implemented without added cost within the power supply by the same line of reasoning. Furthermore, a fully integrated power supply solution will reduce the motherboard part count reducing beat rate and increasing motherboard yield.

3. Interoperability –
The power supply implementation holds an important advantage over the motherboard implementation in the area of interoperability. This specifically pertains to OEMs and resellers who plan to build Instantly Available PCs by assembling components. Placing the circuitry within the power supply provides a unified, fully integrated solution removing the possibility of incompatibility between the switching circuitry on the motherboard and that in the power supply. Motherboard OEMs wishing to achieve Instantly Available features will need only to design their boards according to the current **ACPI specification** (http://www.teleport.com/~acpi) and the Instantly Available PC Design Guide (http://developer.intel.com/design/power/pcpower.htm).

The Ideal Solution – Dual Mode Power Supplies
After taking a closer look, the fully integrated dual mode power supply solution holds many advantages and demonstrates that added features and capabilities do not necessarily equate to absolute increased costs. Although initial power supply cost may increase, eventual volume reduces costs and makes the integrated power supply option more and more attractive. In the end, a fully integrated dual mode power supply is clearly the best long term solution for all OEMs, not to mention a solution that is much more reliable for the end user.

About the Author
Patrick Bohart is a technical marketing engineer in Intel's Platform Marketing group. He focuses on developing technical collateral and information to help accelerate the adoption of Instantly Available PC technology by the PC industry.

For More Information
To stay up to date on Instantly Available Power Managed (IAPM) PC news, visit the **IAPM technology page** in Platform Solutions (http://developer.intel.com/solutions/tech/power.htm).

Read "A New Wakeup Call for the **Instantly Available PC: 3.3 Vaux**," by Gary Solomon, Intel Architect, in Issue #4 of Platform Solutions.
http://test.intel.com/solutions/archive/issue4/stories/top4.htm

Instantly Available PC System **Power Delivery Requirements and Recommendations** (Power Supply '98) Specification (http://developer.intel.com/design/power/supply98.htm).

Technology News:

Technology News—Intel Developer Forum

Tuesday – February 17, 1998

Intel Chairman and CEO Andrew S. Grove Gives Keynote at Intel Developer Forum

Dr. Andrew S. Grove addressed over 1,000 computer industry developers at the Intel Developer Forum on February 17th. Dr. Grove discussed Intel's view that all computing will move into "segments", and the impact of this on Intel's processor and platform products as well as those developing and designing products for Intel Architecture based PCs. He shared his vision for a world in which all computers were connected driving tremendous growth and opportunity for PC developers as IA based PCs move into all segments from entry level, basic computers to enterprise level computing. *Dr. Grove has provided the* ***feature story*** *in this special IDF issue of Platform Solutions based on his keynote address. The* ***webcast archive*** *of his entire keynote speech can also be found at the IDF* (http://developer.intel.com/design/idf/keynote.htm).

http://www.intel.com/pressroom/archive/releases/CN021798.HTM

Intel and Leading Server Vendors Introduce Server Management Hardware Interface IPMI

Intel today announced a new management specification for Intel architecture-based servers intended to enhance server management, drive down the total cost of ownership (TCO) and compliment Intel's Wired for Management initiative. The specification, called the Intelligent Platform Management Interface (IPMI), was announced by Intel, Dell*, Hewlett-Packard Company*, and NEC* to provide a standard interface to hardware used for monitoring a server's physical characteristics such as temperature, voltage, fans, power supplies, and chassis.

http://www.intel.com/pressroom/archive/releases/sp021798.HTM

Intel740 Graphics Chip Receives Eager Industry Support From IHVs and ISVs

The Intel740 was introduced on February 12th as a compelling, next generation 3D graphics accelerator in limited production today (scheduled to be in full production in March 1998) that takes advantage of the AGP Interface Specification. Its unique HyperPipelined 3D architecture and video acceleration provides outstanding graphics realism making visual computing a reality on mainstream PCs. Several hardware vendors showcased their Intel740-related technology today at IDF today.

http://www.intel.com/pressroom/archive/releases/agp21798.HTM

New Server System Infrastructure (SSI) initiative Described at IDF

Intel and leading industry server vendors are working to define open industry specifications for common elements within server systems - the infrastructure. At IDF the draft specification is being discussed with developers. SSI will establish a specification for two common elements of today's and tomorrow's server system infrastructures: power supplies and electronic bays. Visit the new SSI technology page in PSN for more details on the initiative and plans for the specification.

http://developer.intel.com/solutions/tech/ssi.htm

USB Specification Version 1.1 Proposed Modifications Discussed at IDF

The USB core group is deliberating a number of clarifications and enhancements to the current USB Specification version 1.0. An overview of these proposed modifications for incorporation in USB Specification version 1.1 is being discussed at IDF, together with the implications to hardware and software developers. The focus of the update to the specification is to enable more robust designs through examples, guidelines, clarifications and bus usage models. For more details on the proposed modifications visit the USB technology page in Platform Solutions.

http://developer.intel.com/solutions/tech/usb.htm

Wednesday – February 18, 1998

Intel Introduces IOMETER, Server I/O Performance Measurement Tool

Intel announced at IDF the availability of Iometer, a new server I/O performance measurement tool that Intel is distributing at no charge to server system developers at IDF and via the Intel Web site. Iometer provides a simple way to gauge server I/O performance and reliably compare results with those obtained from other server tests, conducted in other labs using Iometer.

http://www.intel.com/pressroom/archive/releases/sp021898.HTM

Intel Releases Platform Performance Tuning and Integration Tools for OEMs & IHVs

Intel announced the availability of a family of platform tools, the Intel Performance Evaluation and Analysis Kit (IPEAK), for PC OEMs and Independent Hardware Vendors (IHVs) to dramatically improve analytical and diagnostic capabilities in platform integration and performance tuning. The tools were developed to assist PC OEMs and IHVs in shortening product time-to-market when adopting new platform technologies and standards, such as AGP, ACPI and Instantly Available PC.

http://www.intel.com/pressroom/archive/releases/cn021898.HTM

Intel & IBM Alliance Breaks New Ground in Systems Management

Demonstrating their commitment to lowering the total cost of ownership of business computing and increasing uptime of PCs in an connected environment, Intel and IBM announced on a new network connection technology called Alert on LAN. This new technology will enable PCs to send immediate alerts to network administrators when there are hardware or operating system failures, or evidence of tampering. The unique value added of Alert on LAN is its ability to generate these alerts, even if the system is powered off and/or the operating system is not yet loaded.

http://www.intel.com/pressroom/archive/releases/WM021898.HTM

IBM & Intel Target DB2 and VI Architecture for New Scalable DB Cluster Solutions

IBM announced its plans to work closely with Intel on the optimization of IBM's DB2 Universal Database on Intel Architecture, providing customers with superior performance, scalability and reliability for mission-critical, line-of-business operations. The joint development effort will focus on the use of the Virtual Interface (VI) Architecture for standards-based clustering.

http://www.software.ibm.com/news/3082.html

Thursday – February 19, 1998

Intel Announces Support for DVD WG-4 Audio Specification for Advanced PC Audio

Intel Corporation today announced its membership into the DVD WG-4 Audio Working Group and support for its proposed DVD audio specification as a next step in bringing next-generation consumer quality audio to the PC platform. Intel is the first computer industry representative invited to join the predominantly consumer electronics forum that has been meeting with music industry representatives for the last two years. Intel supports the proposed WG-4 format as the most cost-effective and PC-friendly and will provide input to the committee in its efforts to make this an accepted standard.

http://www.intel.com/pressroom/archive/releases/dvd21998.HTM

Digital Content Protection Solution Proposed by Leading CE and PC Companies

Intel, Hitachi, Matsushita (Panasonic), Sony and Toshiba announced today a joint proposal to protect digital video and audio content while providing a robust, fast and transparent method for transmitting and receiving digital content between a variety of products such as PCs, high-definition televisions, set-top boxes, digital VCRs and DVD players. The entertainment industry is supportive of this joint proposal. In addition, as a digital solution, the proposed solution also addresses poor quality degradation that currently exists when consumers copy or play back analog video and audio numerous times.

http://www.intel.com/pressroom/archive/releases/con21998.HTM

Quantum Announces Next Generation ATA/IDE, Ultra ATA/66 Disk Drive Interface

Quantum Corporation today announced at IDF the next generation ATA/IDE Ultra ATA/66 disk drive interface, a Quantum patented-technology, which will be supported in future Intel chipsets.
Ultra ATA/66 not only doubles the bandwidth over ATA/33, but also significantly increases the robustness of the interface to the hard drive with improved cabling. Ultra ATA/66 leverages technology from the current Ultra ATA/33 interface, thereby preserving the total cost of ownership for personal computer (PC) users. The interface also enhances data integrity to the ATA interface with improved timing margins and the use of Cyclical Redundancy Check (CRC). CRC is a data protection verification, which helps ensure the integrity of transferred data.

http://www.quantum.com/corporate/pr/ultraata.html

Technology News – Extra

PC 99 System Design Guide 0.3 Now Available For Review

The preview draft 0.3 of the PC 99 System Design Guide is now available for industry review. This guide is intended to advance the quality of PC hardware, firmware, and device drivers for products designed for 1999 production by encouraging PC hardware platform initiatives and technical capabilities resulting in improved user satisfaction and market segment growth. Your complete preview copy can be downloaded from Intel's PC 99 web site and feedback can be sent via email to PC99@intel.com or PC99@microsoft.com

http://developer.intel.com/design/desguide/index.htm

Intel Extends PC Health Management To Mobile PCs with LANDesk® Client Manager v3.2

Intel has introduced LANDesk® Client Manager v3.2, the latest version of its PC health monitoring tool, which helps businesses reduce the costs of business computing. This release extends PC management capabilities from high-performance desktops into the dynamic mobile environment and will make mobile PCs that are not always connected easier to remotely manage, subsequently helping facilitate reduced mobile PC maintenance.

http://www.intel.com/pressroom/archive/releases/LD021098.HTM

New Version Of VTune, Multi-Language Software Optimization Tool For IA Platforms

On February 9th at Software Developer '98 West, Intel announced its newest version of VTune, a visual performance tool for tuning and optimizing software written in C, C++, Fortran, Java* and Visual Basic on Intel Architecture platforms. VTune 3.0 identifies software performance problem "Hot Spots" for most programs in less than 30 seconds. Thousands of developers currently use Intel VTune or have downloaded evaluation copies of VTune from Intel's web site.

http://www.intel.com/pressroom/archive/releases/CN020998.HTM

New Java on Intel Architecture White Paper Available for Download*

Existing performance results show that Intel processors are exceptionally efficient at running Java programs (reference – http://www.webfayre.com/pendragon/jpr –). Intel has conducted a number of studies to determine what else could be done to deliver even greater Java performance. The studies conclude that Java runs exceptionally well on the Intel architecture because the architecture and instruction set were designed to be general purpose so it easily adapts to various computational scenarios. The Intel Architecture (IA) performs well on integer programs as well as on scientific calculations and is well suited to the dynamic nature of Java. Download the complete white paper at the Business Platform page in PSN.

http://developer.intel.com/solutions/platfms/business.htm

Wired For Management Interoperability Workshop Draws 89 New Wfm-Capable Systems

Intel's industry-led efforts to provide easy-to-manage PCs continued gaining momentum at the 3rd Wired for Management (WfM) Interoperability Workshop on January 27th. A total of 17 computer hardware system vendors and six software vendors began final testing of their products' management capabilities and interoperability on a total of 89 platforms. For the first time, mobile PCs and servers were tested, along with desktop PCs, during the event hosted by Intel. Systems tested have been designed according to the Wired for Management (WfM) Baseline Specification.

(http://www.intel.com/pressroom/archive/releases/WM012798.HTM)

Industry Events:

Intel Developer Forum
February 17-19, San Jose, CA, USA

Go *"Beyond the Spec"* at this second bi-annual hardware developer event hosted by Intel, at the San Jose Convention Center. The Intel Developer Forum (IDF) is the industry's premier event for hardware developers. Get implementation tools, detailed training and knowledge, directly from Intel's chief technology architects, on the latest technologies driving the hardware platform. IDF covers today's implementation details and tomorrow's technology roadmaps, to help speed the development of new products integrating the current advancements in desktop, mobile, workstation and server platform technologies. In addition, IDF provides a valuable opportunity to establish and strengthen personal working relationships with technology leaders from throughout the PC industry, and from around the world.
For more details on the IDF technical tracks and registration, please visit the **IDF web site** (http://developer.intel.com/design/idf)

Intel Developer Forum DVD Plugfest
February 24-26, Milpitas, CA, USA

This "Plugfest" event, held at the Milpitas Embassy Suites the week after the Intel Developer Forum, will involve interoperability testing of various components of DVD playback for personal computers. Vendors of DVD drives, graphics cards, MPEG hard and soft decoders, OEMs of personal computers, and software content providers will convene to test their products for compatibility over a 3 day period. Stay tuned to Platform Solutions for more information on this event in the future. Register for this event today by visiting the **DVD Plugfest web site** (http://industry-enabling.org/plugfest.htm)

Spring Internet World 98
March 11-13, Los Angeles, CA

Internet World is the industry's oldest and largest Internet, Intranet and World Wide Web event. Internet World events include all categories of products and services within the Internet industry and attracts all segments of the computer industry. For more details and registration, please visit the **Internet World web site** (http://events.internet.com/spring98/spring98.html)

CeBIT
March 18-25, Hannover, Germany

Europe's premier computer, telecommunications, and Information Technology industry fair.
For more details, please visit the **CeBIT web site** (http://www.messe.de/cb98/index_e.html)

WinHEC
March 25-27, Orlando, FL, USA

Windows* Hardware Engineering Conference to be held at the Orange County Convention Center. PC industry event for manufacturers and suppliers of hardware products supporting Microsoft* Windows* family of operating systems. WinHEC brings together technical managers and product developers to examine new technologies for designing future Windows based computers.
For more details and registration information, please visit the **WinHEC web site** (http://www.microsoft.com/hwdev/winhec.htm)

Comdex® Spring '98
April 20-23, Chicago, IL

Key international event for the entire computer industry. For more information please visit the **Comdex Web site** (http://www.comdex.com)

USB Compliance Workshop
April 29, 30 and May 1st, Milpitas, CA

For Invitation, Registration and Details, please visit the **USB Implementers Forum web site**
http://www.usb.org/developers/index.shtml

PCI Expo '98/PC Developers Conference
May 18-22, San Jose, California

PC and industrial developers' conference for hardware and software engineers. Intel will provide kenotes on connectivity and USB.
http://www.annabooks.com

USB Plus Developer's Conference
May 20-22, 1998, San Jose, CA

ATTENTION! The upcoming USB Plus Developer's Conference will be co-produced by USB-IF and Annabooks. In this 3-day conference, you will hear first hand from knowledgeable speakers currently defining and developing state-of-the-art USB products. Registration and program coordination will be handled by Annabooks. To register simply visit the **Annabooks web site**
www.annabooks.com/confer/pciregis.htm

Exhibiting opportunities are available: Contact Shauna Wilson via email--- dwilson1@san.rr.com or call 619-689-4942.

Intel Networking Events & Training
For Intel's events and training programs on networking products and technologies, please visit the **Intel networking events page** (http://www.intel.com/network/events/index.htm)

- end of Platform Solutions newsletter Issue 6 -